Simulating the Impact of Hybrid Functionality on CHAPS Banks

A. Kemal Ercevik
and
John P. Jackson

Working Paper no. [###]

DRAFT: Please do not quote without the written consent of the authors
(This version: 31st July 2007)

Financial Resilience Division, Financial Stability, Bank of England, Threadneedle Street, London, EC2R 8AH.

kemal.ercevik@bankofengland.co.uk
john.jackson@bankofengland.co.uk

The views in this paper are those of the authors and do not necessarily reflect those of the Bank of England. Helpful comments and advice were received from a number of parties during the preparation of this work. In particular, the authors are very grateful to Neville Arjani, Mark Manning, Andrew Mason, Kimmo Soramaki, participants at the Bank of England's February 2007 practitioners seminar and June 2007 Economics of Payments workshop, and attendees at a central bank policy workshop at the BIS in March 2007, for their helpful comments. Special thanks also go to Joanna McLafferty for her dedicated research assistance and to Matti Hellqvist for his technical assistance with the Bank of Finland simulator. This paper was finalised on [dd/mm/yyyy].

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at www.bankofengland.co.uk/publications/workingpapers/index.html.

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone +44 (0)20 7601 4030, fax +44 (0)20 7601 3298, email mapublications@bankofengland.co.uk.

ISSN [##############] (on-line)

Contents

Abstract

This paper uses a simulation methodology and real payment data to quantify the impact of introducing a centralised receipt-reactive queue on the liquidity demands faced by CHAPS Sterling banks. Significant liquidity savings are achievable if banks choose to submit a very high proportion of payments (>90% by value) into the queue. The relationship between values queued and savings achieved is shown to be non-linear. Only limited delays in settlement are seen following the introduction of queuing. Liquidity savings are distributed unevenly, with the largest CHAPS banks seeing no benefits while the smaller CHAPS banks see very large savings. A synthetic payments dataset is used to demonstrate that the key aggregate and bank level impacts of queuing functionality found using CHAPS data hold more generally. The key determinant of the impact of hybrid functionality is shown to be the liquidity recycling ratio being achieved in the existing system, which in turn is influenced by the number of direct participants using the system and the volumes of payments they process. Comparing findings with real and synthetic data suggests that CHAPS banks already adopt payment submission strategies that significantly reduce the volatility of their liquidity demands.

Key words: payment system simulation; hybrid functionality.

1. Introduction

This paper uses real payment data to test the possibility that users of the CHAPS Sterling system could benefit from the introduction of a hybrid payment system design. Simulations are used to quantify the impact of such a design change on users in terms of the impact on both their liquidity usage and the degree of settlement delay introduced. The analysis is extended by generating synthetic payment data and using it to probe the extent to which our findings are CHAPS specific.

CHAPS Sterling is the UK's large value payment system (LVPS) with a mean daily turnover of £231 billion.[1] It operates on a real-time gross settlement (RTGS) basis: payments are processed transaction by transaction with simultaneous debit of the payer and credit of the receiver in accounts held at the Bank of England. RTGS eliminates interbank credit risk by providing immediate finality of payments: once a transaction has settled it is irrevocable and cannot legally be unwound (for a detailed explanation of RTGS systems, see BIS (1997)). This important feature of RTGS has led to its adoption in LVPSs worldwide, as banks and central banks have found the settlement risk inherent in deferred net settlement (DNS) systems to be too high given the size of interbank exposures that can occur. Although RTGS in central bank money eliminates settlement risk, it can have the undesired consequence of increasing the cost of making payments and potentially increasing liquidity risk faced by banks. When payments are settled gross rather than being netted out at the end of the day, banks typically require more liquidity to make their payments. CHAPS banks, for example, require on average three times more liquidity under RTGS than they would have needed under a DNS system with multilateral end-of-day netting.[2]

The past decade has seen a growing trend of adoption of hybrid payment system designs for the settlement of large-value payments in developed countries.[3] Hybrid systems seek to be liquidity-efficient without introducing significant amounts of settlement risk by combining features of both RTGS and DNS systems. A defining feature of these systems is that payments can be centrally queued, with their release conditional on certain criteria, such as the arrival of offsetting payments. Offsetting occurs when two or more payments are settled simultaneously. Although in legal terms settlement is still gross (i.e. each transaction is settled with finality individually) payments that are offset can be thought of as having generated their own liquidity, as offsetting has the same economic effect as the netting of payments.

Two main types of hybrid systems have emerged to date. One type, which has evolved from DNS, is called continuous net settlement (CNS). Examples of such systems include the PNS in France and CHIPS in the US. Although settlement risk is significantly reduced in CNS systems (payments that are offset and settled with finality in batches intraday are not dependent on the subsequent settlement of other payments in the system), it is not completely eliminated. The second type of hybrid system incorporates a queuing facility into RTGS, in effect creating multiple streams into which banks can channel their payments: typically a time-critical (RTGS)

---

[1] Payment Systems Oversight Report 2006 (www.bankofengland.co.uk/publications/psor/psor2006.pdf).
[2] Source: Bank of England payments database and authors' calculation for December 2006.
[3] Calculations in unpublished work by Bech *et al.* (2007) show that in 1999 only 3% of wholesale payments by value in CPSS countries were settled over payment systems with a hybrid design, while in 2005 that had risen to 32%.

stream and an offsetting (queuing) stream. Examples of systems with this type of functionality are RTGS^{plus} in Germany and TARGET2, the centralised LVPS of the Eurosystem due to go live later this year. Such queue-augmented RTGS systems do not give rise to settlement risk[4], as all queued payments are legally considered to be final at the time they are offset. For a more detailed explanation of hybrid system designs see McAndrews and Trundle (2001) and BIS (2005).

Although queue-augmented RTGS systems can help reduce liquidity risk without introducing settlement risk, they can introduce settlement delays which may impose additional costs on banks. A number of theoretical and empirical papers assess this fundamental trade-off between liquidity efficiency and settlement delay in an attempt to find the optimal settlement arrangement for LVPSs. Johnson, McAndrews and Soramaki (2004) use simulation techniques to assess the impact of various types of hybrid functionality on Fedwire, the RTGS system used to settle large value US dollar payments. The paper finds that one mechanism in particular, receipt-reactive gross settlement (RRGS), can potentially reduce participants' costs of obtaining intraday credit, whilst only modestly delaying the average time of settlement.[5] RRGS is a novel queue release mechanism proposed by the paper that conditions the settlement of queued payments on the arrival of incoming payments. This feature ensures that all the liquidity posted by a bank is reserved solely for making time-critical payments. The paper recognises that the introduction of RRGS functionality would provide a good incentive for banks to submit payments earlier in the day, but does not attempt to incorporate this behaviour in its simulations.

Willison (2004) and Martin and McAndrews (2007) both use game-theoretic models to predict and compare equilibria for RTGS and hybrid system designs. Willison defines the first-best solution in terms of the trade-off between cost of liquidity and operational risk caused by payment delay, and finds that the first-best is unattainable under RTGS. He also finds that a hybrid payment system design outperforms RTGS when payments can be offset either in the morning or all day. Martin and McAndrews find that a balance-reactive hybrid system (a system where settlement of queued payments is conditional on a participant's account balance) can provide higher or lower welfare than RTGS depending on certain criteria, such as the cost of delaying payments and the proportion of time-critical payments in the system. Welfare is defined in terms of the cost of liquidity and cost of delay borne by the participants. A receipt-reactive system on the other hand weakly dominates RTGS: it can achieve a level of welfare at least as high as, if not better than, RTGS.

The purpose of this paper is to investigate the possibility of achieving liquidity savings in CHAPS by simulating the impact of hybrid functionality[6] using historical data. We attempt to fill some of the existing gaps in the literature outlined above by:

---

[4] As long as receiving banks do not anticipate payments they are due to receive in the queue and credit beneficiaries' accounts before the payment has been settled / offset with finality. To avoid this, some hybrid systems provide very limited information about payments in the central queue (payer, payee, value only) until finality has been achieved.

[5] This cost reduction is based on the Federal Reserve's method of charging for intraday credit, which is a fee based on banks' average overdrafts calculated at the end of each calendar minute, and does not necessarily apply to a system where intraday credit is free but collateralised, such as CHAPS.

[6] CHAPS is not considered to be a hybrid system, although it does feature a central queue. A 'circles' process is run once a day, and can be run additionally by the RTGS system controller at any time during the day in order to resolve

- Analysing the aggregate impact of hybrid functionality on CHAPS banks in terms of costs of liquidity and degree of delay introduced by RRGS[7], experimenting with different criteria for time-criticality. We find that significant liquidity savings are achievable where banks choose to submit a very high proportion of payments (>90%) into the receipt-reactive queue, and show that the relationship between values queued and savings achieved is non-linear. The level of settlement delay introduced also increases in a non-linear fashion as queued values rise, but does not reach excessive levels under any assumptions used.

- Assessing the impact of hybrid functionality at the individual bank level. We show that liquidity savings are unevenly distributed across banks, with an inverse relationship between the value of payments made by a bank and the liquidity saving benefits to that bank from the introduction of RRGS. For the largest CHAPS banks liquidity demands can even rise slightly under RRGS.

- Endogenising banks' likely payment submission behaviour under RRGS by submitting some non time-critical payments earlier in the day. We show that where significant levels of payment delay exist in an RTGS system, the value-weighted average time of settlement can be brought forward significantly by the introduction of receipt-reactive functionality due to the impact on payment behaviour. We do not expect this effect to alter our CHAPS results significantly as payment delay is not a significant problem in the system and some of the delay that does exist is introduced by use of bilateral limits that would likely remain under RRGS.

- Generating synthetic payment data to probe the extent to which our findings are CHAPS specific and investigate the key determinants of the impact of hybrid functionality. We find that the impact of hybrid functionality is strongly influenced by the liquidity recycling ratio being achieved under RTGS, which in turn is influenced by the number of direct participants in the system and the amount of payments they process. We corroborate that key aggregate and bank level findings for CHAPS hold more generally. The differences between real and synthetic data in the impact of RRGS on volatility of liquidity demands suggest that CHAPS banks already submit payments using strategies that reduce the volatility of their liquidity needs.

The rest of this paper is organised as follows. In the next section, we explain our simulation methodology by defining RRGS and introducing the metrics we use to measure the impact of hybrid functionality on liquidity usage and payment delay. In Section 3, we report results of our simulations, interpreting them in Section 4 and examining possible policy implications. We conclude in Section 5 by summarising our key results and outlining some possible extensions to our work. More detailed information on how we generate our synthetic payments datasets can be found in the Annex.

---

gridlock situations when participants have insufficient funds on their accounts, which momentarily locks payments waiting in the central queue and multilaterally attempts to offset as many of them as possible. This is not, however, used as a liquidity saving feature in its current form; banks prefer to queue payments in their internal schedulers rather than submit them to the central queue. A gridlock has never occurred to date due to the posting of ample amounts of liquidity by CHAPS banks at the start of each day.

[7] We would also have liked to simulate a *balance-reactive* hybrid system design, similar to the one being developed for TARGET2 or the one already being used in the RTGS^{plus} system. Unfortunately, such a queuing algorithm is not available in the current version of the Bank of Finland simulator (BoF-PSS2 v2.2.5). Simulating this functionality would be an obvious extension to our work, once the functionality becomes available.

2. Method

We run simulations to test the possibility of making liquidity savings in CHAPS by complementing the existing RTGS stream with RRGS, using a month of historical payments data.[8] [9] We use the Bank of Finland payment and settlement simulator (BoF-PSS2), which is described in detail in Leinonen and Soramaki (2003). Section 2.1 describes the RRGS functionality in more detail. Time-critical payments and the metrics used to measure the impact of our simulations are defined in sections 2.2 and 2.3 respectively. Sections 2.4 and 2.5 explain how we endogenise banks' payment submission behaviour and generate synthetic payments data.

2.1 Receipt-reactive gross settlement

The RRGS algorithm can be viewed as an extreme case of liquidity reservation functionality. All the liquidity posted by a bank into the system is reserved to allow payments which that bank designates as high priority to be settled immediately. The same bank's low priority payment messages are released for settlement, on a first-in first-out (FIFO) basis, only where they can be settled using liquidity from the arrival of incoming funds within a pre-specified period of time. Johnson *et al.* (2004) use calendar minutes as the time intervals in their paper: in any minute the algorithm allows the release of as many payments from the front of the queue as is possible to offset, but not exceed, the amount of incoming funds received in that minute. They use this approach because in Fedwire charges for banks' daylight overdrafts are calculated at the end of each calendar minute.

In our simulations it is appropriate to take the entire CHAPS day as one continuous period[10], as the key determinant of the cost of posting liquidity in CHAPS is the *maximum* liquidity needed throughout the day. This means that the RRGS algorithm runs on a continuous basis: a payment received at 9am can cause the release of a payment entered into the queue at 10am or even 4pm as long as aggregate payments received by bank i >= aggregate *queued* payments sent by bank i at that point in time, including the queued payment(s) being released. Under RRGS, a bank does not necessarily have to post liquidity to cover the gross value of *all* its outgoing time-critical payments: incoming funds can be used to finance both time-critical and non-time critical payments. The only distinction is that a time-critical payment will never queue, so if a bank has

[8] December 2006 – a very 'clean' month with no CHAPS settlement extensions. 19 working days in total.
[9] In addition to the RRGS algorithm, the Bank of Finland offers a range of bilateral and multilateral offsetting algorithms whose impact on CHAPS could in principal be simulated. In practice however this is not a viable alternative as it is not possible to simulate the impact of such an algorithm while continuing to allow time-critical payments to be settled immediately without carrying out two distinct simulations: a pure RTGS simulation for time-critical payments, and a continuous net settlement (CNS) simulation with multilateral offsetting for the remaining transactions. Simulating such an arrangement using CHAPS data results in banks requiring *more* liquidity than where all the payments are settled RTGS (i.e. without offsetting) under one simulation. This occurs because liquidity recycling is disrupted by the 'splitting' of transactions between two accounts. A surplus of liquidity in one account cannot be used to fund a deficit in the other, requiring further liquidity injection, which would not have been the case if the funds had been posted onto a *single* account (as is the case, for example, with the RRGS algorithm which has multiple streams but uses a single account).
[10] We also investigate whether adopting a greater number of distinct periods over which incoming payments are cumulated impacts our results by experimenting with 1 hour and 5 minute intervals. We find that estimated collateral posting by banks is not significantly affected, as mean liquidity requirements across the month slightly decrease, but this is offset by an increase in volatility (standard deviation) across the 19 days: for a detailed explanation of the metrics used to quantify the liquidity burden faced by banks and their impact on estimated collateral postings see Sections 2.3.1 and 4.4 respectively.

not received enough payments to make a RTGS payment, as reflected by its account balance, it will have to post additional liquidity to cover that payment.

Finally, it is worth noting that although we refer to the term offsetting throughout the paper, this misleadingly suggests that offset payments are always released for settlement at the same instant. This can and often does happen, but there are also instances where a payment entered into the queue is conditionally released in response to a payment received much earlier in the day. RRGS should therefore be viewed as a mechanism for co-ordinating the immediate settlement of time-critical payments, together with the conditional release of the remaining payments against incoming payments: where we use the term *offsetting* in the context of RRGS, we specifically mean offsetting through *conditional release*. Also, if there are any unsettled payments remaining in the queue at the end of the day, we assume these are settled on a multilateral net basis. The chart below illustrates how the RRGS mechanism works:

Chart 1: Dynamics of a bank's balance under RRGS

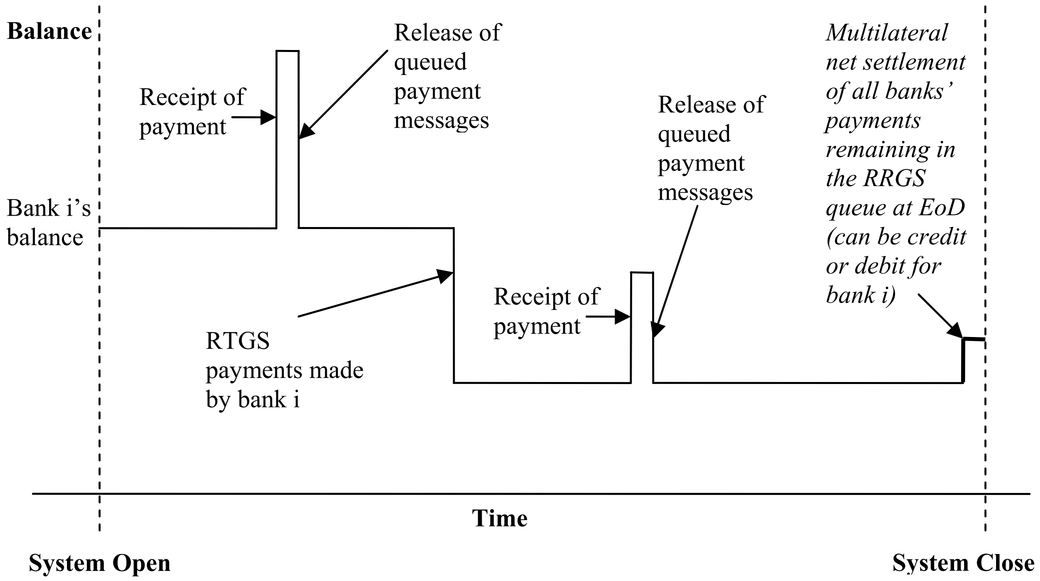

Source: Bank of Finland PSS2 User Manual v 2.2.0

2.2 Time-criticality

Throughout our analysis we always require time-critical payments[11] to be settled immediately but allow the remaining payments to be queued, waiting for incoming payments to trigger their release. Anecdotal evidence suggests the main payment types of payments which are considered by CHAPS banks to be truly time-critical are payments that are either: (a) to or from the account of CLS[12] [13]; (b) extremely large in value; or (c) require prompt settlement by customers e.g.

---

[11] The term time-critical payment refers to a payment for which the sender deems there are significant private or social costs to delay; indeed, in some cases, failure to pay at, or by, a given time intraday may constitute technical default.
[12] CLS is the Continuous Linked Settlement system which provides payment versus payment settlement of FX transactions, see www.cls-services.com for more details. Users' sterling pay-ins to CLS are made across CHAPS.
[13] More generally, pay-ins to all ancillary systems are typically viewed as highly time-critical. However, CLS is the only payment of this type which is made across CHAPS. Settlement of BACS, the UK retail clearing system, takes

house purchases. It is not possible to identify the last category of time-critical payments in our dataset but we always treat payments to and from CLS Bank's CHAPS account as high priority and settle them RTGS. As we have no other obvious way of identifying precisely other categories of time-critical payments in our data, we experiment with two approaches to proxy time-criticality:

(i) Payments of a size greater than or equal to a certain value threshold are time-critical and have to be settled immediately; others only have to be settled sometime during that day (at the latest by the end of day) and could therefore be queued and offset. We experiment with different value thresholds by treating payments larger than £100m, £500m and £1bn as time-critical.

(ii) A random percentage (*r*) of payments are time-critical and must settle RTGS. This approach is used in Johnson *et al.* (2004). For ease of comparison we adopt values of *r* to ensure that the same proportion of payments by value is treated as time-critical under approaches (i) and (ii).

In practice, banks' time-critical payments are likely to be somewhere in between the two sets of payments we identify under the approaches above. Our method therefore enables us to report a range of possible liquidity saving / settlement delay trade-offs, from which both extremes of the impact of RRGS on CHAPS banks can be observed.

2.3 Metrics for comparison

Throughout the paper we compare metrics from our simulations with those obtained from a pure RTGS simulation. We focus on the following measurements to quantify the impact of RRGS on the trade-off banks face between liquidity usage and settlement delay.

2.3.1 Metrics used to quantify the liquidity burden faced by banks

We use various measures to identify relevant indicators of banks' liquidity posting requirements. One simple measure is the mean of the maximum daily net debit position faced by each bank across the month, which is then summed across banks to get an aggregate figure for the system. In practice it would not be feasible for banks to post this exact amount into the system *ex ante* as it is determined by the submission behaviour of all banks in the system, but it gives a good indicator of liquidity demands faced by banks. One potential drawback of this metric is that it cannot quantify changes in the volatility of liquidity demands that may be caused by RRGS. To capture this we quantify how the standard deviation of maximum daily net debit positions over the month changes for each bank. We also display results for the liquidity required by each bank to cover the single maximum net debit position faced across the entire month.

---

place as a non-CHAPS transfer across banks' RTGS accounts and is not included in our dataset. Similarly, liquidity transfers to CREST for settlement of DvP transactions occur across other liquidity transfer accounts.

Mean of maximum liquidity requirement across the month for bank i:

$$\overline{MLR}_i = \frac{1}{n}\sum_{j=1}^{n} MLR_{ij}$$

where:
n = number of CHAPS days being simulated
MLR_{ij} *= maximum liquidity requirement of bank i on day j*

Standard deviation of maximum liquidity requirement across the month for bank i:

$$Sd_i = \sqrt{\sum_{j=1}^{n} (MLR_{ij} - \overline{MLR}_i)^2 /(n-1)}$$

Maximum of maximum liquidity requirement across the month for bank i:

$$Max_i = \max(MLR_{ij})$$

2.3.2 Metric used to quantify average settlement time

The basic indicator that we consider to quantify the average settlement time under different simulations is the value-weighted average settlement time (*AST*), calculated as follows:

$$AST = \frac{\sum_i t_{s,i} a_i}{\sum_i a_i}$$

where:
$t_{s,i}$ *= settlement time of payment i*
a_i *= value of payment i*

Early settlement is desirable from an operational risk perspective: the earlier payments are settled, the lower the risk of having large amounts of payments remaining to be settled following an unexpected operational outage.

2.3.3 Recycling ratio

We use the recycling ratio, *rr*, to calculate the liquidity efficiency of the system before and after the introduction of RRGS. This is based on the method used in Becher *et al* (2007) and is calculated as follows:

$$rr = \frac{\sum_{j=1}^{n} \sum_i a_{ij}}{\sum_{j=1}^{n} \sum_i MLR_{ij}}$$

The ratio measures how many times the system can recycle £1 of liquidity by comparing the total value of payments submitted with the maximum value of liquidity that needs to be posted by members of the system to allow settlement to occur.

2.4 Endogenising payment submission behaviour

Anecdotal evidence obtained from dialogue with CHAPS users explains how banks might alter their behaviour in response to a liquidity saving mechanism such as RRGS. Banks would have an incentive to submit payments to the central scheduler as early as possible to maximise the benefits of offsetting, safe in the knowledge that their liquidity posted at the start of day would be reserved solely for making time-critical payments. This incentive to submit payments earlier in the day is also predicted by recent literature (in Johnson *et al* (2004) and implicitly in Willison (2004)).

We test the impact of this prediction by artificially changing the submission times of payments in some of our simulations. However, banks have made it clear to us that their ability to submit payments earlier in the day would depend on having prior knowledge of individual payments, which varies from bank to bank according to the mix of payments they process. We simulate the impact of these behavioural changes by submitting randomly selected 20, 50 and 80% of non-time critical payments greater than £1m at the start of day instead of their original submission times.[14] In doing this, we incorporate findings on overnight loans data from Bank research currently underway by bringing forward all loan repayments to the beginning of the day (as these are known in advance) and none of the new loans being generated on that day (as these are not known until later in the day).[15]

2.5 Synthetic payments

We generate datasets of synthetic payments to test the extent to which our findings are CHAPS-specific (for detailed methodology, see Annex 1). Using such data allows us to vary the numbers of banks in a system, the numbers of payments being settled, and the distribution of values and volumes of payments across banks to see what drives the impact of RRGS.

---

[14] Banks have told us that they tend to make their small value payments (we assume this means for value <£1m) immediately as they are received without queuing them in their internal schedulers, reflecting the fact these payments do not impose significant liquidity demands. It would not therefore be realistic to alter the submission times for these payments as they are not subject to delay. Similarly, we do not change the submission times for time-critical payments, since by definition these could not have been delayed.

[15] We identify overnight loan payments in CHAPS using a Matlab program which matches all the overnight loans in our dataset to loan repayments the following day. The program uses a similar methodology to that outlined in Millard & Polenghi (2004).

3. Results

We display system-wide results using CHAPS data in Section 3.1 (aggregate individual banks excluding Bank of England and CLS Bank[16]) before looking at the impact of RRGS on individual CHAPS banks in Section 3.2. We then compare our findings with results obtained using a dataset of synthetic payments in Section 3.3.

3.1 Aggregate findings

Table A overleaf shows that no significant liquidity savings or settlement delays are observed where only half of all payments by value are submitted to the receipt-reactive stream. Significant mean liquidity savings start to occur as the proportion of queued payments is raised; at the same time minor increases in settlement delay can be observed. The impact of both effects is increasing in the numbers of payments queued.

Mean liquidity savings are not influenced by the method through which time-critical payments are selected. By contrast, the volatility (standard deviation) of liquidity demands is affected. Volatility is unchanged or even increases slightly under RRGS where a value threshold is used, while volatility falls where a volume-based selection method is adopted. This difference in volatility is also evidenced by a corresponding difference in liquidity savings based on the maximum liquidity requirement measure.

Earlier submission of a subset[17] of non-time critical payments has little impact on observed liquidity savings under RRGS, but can significantly bring forward the average time of settlement when compared to RTGS.

---

[16] We treat RBS and NatWest as a single entity throughout our simulations, even though they have separate CHAPS Sterling settlement accounts. We therefore report findings for only 12 banks in our results even though there are 15 direct CHAPS Sterling participants in our dataset.
[17] Based on certain criteria – see Section 2.4 for a detailed explanation.

Table A: Impact of RRGS on CHAPS at the aggregate level

Time-critical payments[a]		Early submission[c]	% Δ Liquidity requirement			Δ Settlement time	
Criteria	Proportion[b]		Mean	St dev	Max	AST[d] hh:mm	St dev[e] hh:mm
≥ £100mn	54%	-	-2	+1	0	+00:01	00:00
≥ £500mn	12%	-	-10	+10	+3	+00:12	-00:01
≥ £1bn	4%	-	-38	0	-9	+00:37	+00:09
Random 50%	51%	-	-1	-2	-1	+00:01	00:00
Random 10%	11%	-	-12	-10	-8	+00:11	-00:01
Random 3%	4%	-	-37	-18	-16	+00:25	+00:04
≥ £100mn	54%	Random 50%	0	+5	+1	-00:57	+00:01
≥ £500mn	12%	Random 50%	-15	+9	0	-01:20	+00:06
≥ £1bn	4%	Random 50%	-37	+5	-4	-00:33	+00:30
Random 10%	11%	Random 50%	-14	-12	-5	-01:15	+00:02
Random 3%	4%	Random 50%	-39	-18	-16	-01:01	+00:12
≥ £1bn	4%	Random 20%	-37	+2	-8	+00:05	+00:15
≥ £1bn	4%	Random 80%	-37	+2	-7	-01:01	+00:43

[a] *CLS payments always treated as time-critical.*

[b] *Proportion of all payments that are time-critical (by value).*

[c] *A dash indicates that original (RTGS) payment submission times have not been altered. In subsequent simulations we extend our analysis by endogenising banks' payment submission behaviour: banks submit all overnight loan repayments and a randomly selected % of payments (which are ≥ £1mn and non-time critical in both cases) to the RRGS queue at the start of day. Submission times of new overnight loans are left unchanged.*

[d] *Change in value-weighted Average Settlement Time (AST) compared to RTGS AST of 11:37am.*

[e] *Change in standard deviation of settlement time across the month compared to RTGS.*

3.2 Bank level findings

Table B overleaf shows disaggregated results for value (≥£1bn) and volume (random 3%) based time-criticality thresholds, grouping banks by the mean liquidity savings they experience when using RRGS. We see that the mean liquidity requirements of the largest two banks increase slightly as a result of the introduction of RRGS. Between them these banks settle half the value of both total payments and of time-critical payments. They both have high liquidity recycling ratios under RTGS.

We see another small group of medium sized banks settling an average of 10% of value and of time-critical payments, which experience moderate mean liquidity savings. Their liquidity recycling ratios prior to RRGS are also moderately high.

Finally, the majority of banks are much smaller in terms of total value and proportion of time-critical payments settled (around 3-4%). They typically have very low recycling ratios under RTGS, and benefit the most from the introduction of an RRGS stream. This pattern of mean liquidity savings being broadly inversely correlated with recycling ratios under RTGS and the sizes of banks' payment flows, is observed for both value and volume based time-criticality thresholds.

Table B: Impact of RRGS on CHAPS at grouped bank level					
Δ Mean liquidity requirement	# Of banks	Avg value settled	Avg proportion of time-critical payments settled	Recycling ratio	
				RTGS	RRGS
≥£1bn time-critical:					
$MLR_i > 0\%$	2	27%	25%	30	28
$0\% > MLR_i > -40\%$	2	9%	13%	14	21
$MLR_i < -40\%$	8	3%	3%	9	24
Random 3% time-critical:					
$MLR_i > 0\%$	2	27%	30%	30	26
$0\% > MLR_i > -40\%$	1	10%	8%	17	25
$MLR_i < -40\%$	9	4%	4%	9	23

3.3 Results using a synthetic payments dataset

Similar simulations were run using our synthetic dataset to attempt to replicate the results found using CHAPS data (i.e. those displayed in the first six rows of Table A, and in Table B). Baseline results are produced for simulations with 10 banks, with the value profile of payment flows and the distribution of payment volume across banks both being drawn from log normal distributions, and with a process of squaring-off of balances taking place during the second half of the payment day. Sensitivity analysis is carried out to examine how the impact of RRGS is altered where the number of banks making payments is increased (we simulate systems with 100 and 1000 banks).

Our results can be found in Table C overleaf. We observe some clear similarities between these results and those displayed in Table A. In particular, no significant liquidity savings are observed where only half of all payments are put into the receipt-reactive stream, but significant mean savings do occur as the proportion of queued payments is raised, with the effect increasing in the value of payments queued. Mean savings are of the same magnitude under value and volume based thresholds for time-critical payments.

One key difference is observed. With the synthetic payments dataset, the volatility of liquidity demands decreases by a greater amount than the decrease in mean liquidity demands for both time-criticality thresholds. By contrast, volatility falls by less than the mean (and in the case of value based thresholds doesn't fall at all) with our CHAPS dataset.

Our sensitivity analysis shows that there is a strong correlation between the number of banks in the payment system and liquidity savings seen under RRGS, with the largest savings observed in our dataset with 1000 banks. This seems to be linked to the observation that recycling ratios fall as the numbers of banks in the system is increased.[18]

We also try a 10 bank simulation where we reduce the overall number of payments from 20,000 per day to 5,000. This halves the system's recycling ratio under RTGS (from 47 to 23), possibly

[18] Simulating a *uniform* distribution of bank sizes and payment values has a similar effect to increasing the number of banks: the less concentrated payments are among a few participants, the more opportunity there is for liquidity savings to be made from co-ordinated settlement of non-urgent payments.

due to the probability of recycling payments falling when there are fewer payments to be made in the system. Correspondingly, the benefit of RRGS doubles on the mean liquidity requirement measure (from 17% to 37%), and the recycling ratio of the system improves by 50% after RRGS (from 23 to 36).

Table C: Impact of RRGS using artificial payments datasets
Log-normal distribution of payment values and bank sizes, 20 days, 400,000 payments

Time-critical payments	Number of banks	Recycling ratio		% Δ Liquidity requirement		
		RTGS	RRGS	Mean	St dev	Max
Largest 50%	10	47	47	0	0	0
Largest 10%	10	47	50	-7	-7	-7
Largest 5%	10	47	56	-17	-14	-17
Random 50%	10	47	47	-1	-2	-1
Random 10%	10	47	50	-8	-14	-14
Random 5%	10	47	58	-19	-28	-22
Largest 5%	100	20	27	-28	-33	-29
Random 5%	100	20	28	-30	-50	-38
Largest 5%	1000	11	20	-45	-44	-40
Random 5%	1000	11	21	-46	-68	-59
100,000 payments						
Largest 5%	10	23	36	-37	-30	-33

Chart 1 below shows how the value of payments submitted is distributed across the dataset of 1000 banks. Chart 2 illustrates how individual recycling ratios are linked to a bank's size. Both charts display the y-axis on a log scale.

Chart 1: Average value submitted (1000 banks)
Bank sizes and payment values log-normally distributed

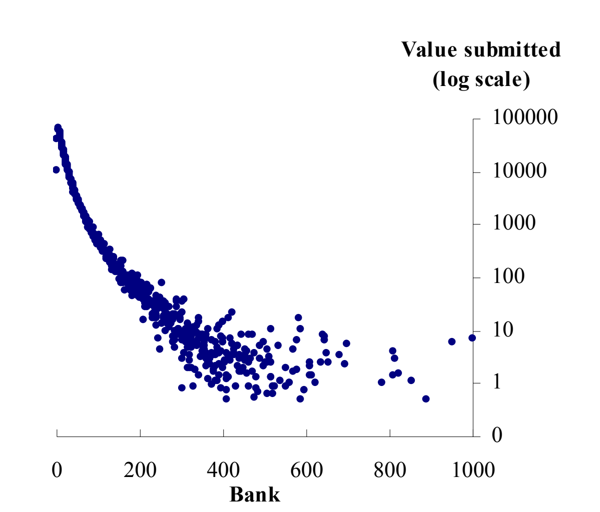

Chart 2: RTGS recycling ratios (1000 banks)
Bank sizes and payment values log-normally distributed

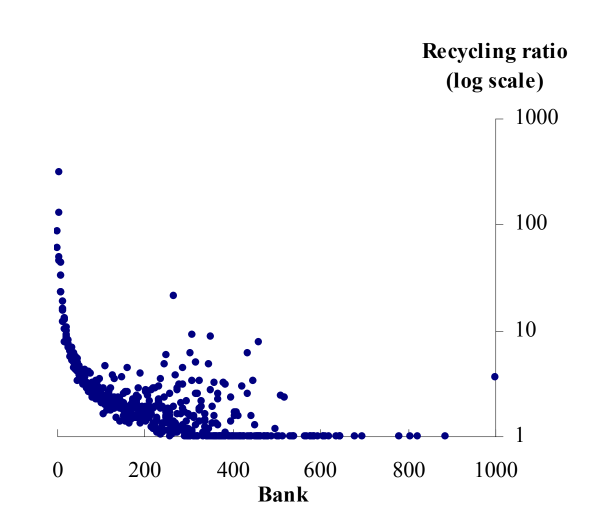

Chart 3 overleaf shows mean liquidity savings under RRGS where the largest 5% of payments are treated as time-critical, by averaging banks in 'buckets' of 50. Although not directly observable in the chart due to averaging, those banks with the highest recycling ratios under RTGS see their liquidity demands increase under RRGS, the same pattern observed in our CHAPS data in Table B. Again we see smaller banks with lower recycling ratios under RTGS making larger savings. The one caveat to this pattern is that the very smallest banks in the sample often do not see any savings as their volumes of payments are not sufficiently large for RRGS to give any benefit (i.e. on a typical day they do not have any offsetting benefits to take advantage of).

Chart 3: Mean liquidity savings under RRGS
Bank sizes and payment values log-normally distributed - largest 5% of payments time-critical

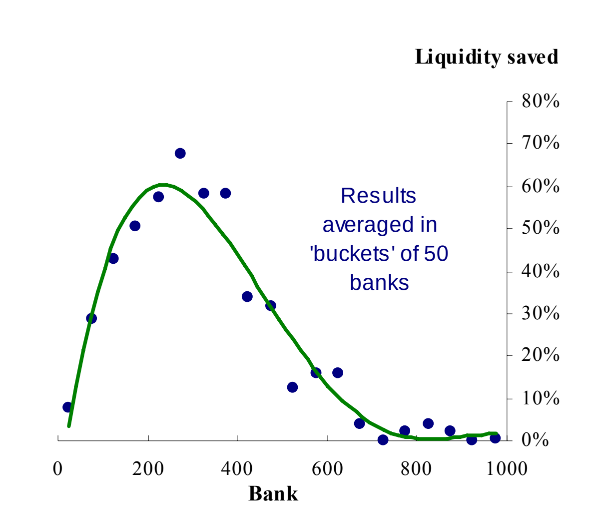

4. Interpretation

This section discusses some of the key findings reported in the previous section. Section 4.1 discusses aggregate findings from CHAPS data. The non-linear increase in aggregate liquidity savings seen as more payments are queued is examined in more detail. In addition the differing impact of volume and value based time-criticality thresholds are discussed. Section 4.2 examines why significant differences in liquidity savings are observed across CHAPS banks. In section 4.3 the bias that might be introduced because our simulations cannot capture the use of bilateral monitoring by CHAPS banks is discussed, with particular reference to its impact on the average time of settlement measure. Section 4.4 considers how liquidity savings observed using CHAPS data might be translated into cost savings, using regression analysis borrowed from James and Willison (2004). Section 4.5 identifies general policy implications of our findings, drawing on similarities and differences observed in the results obtained between the real and generated payments datasets.

4.1 Aggregate findings

Chart 4: Non-linear profile of mean liquidity savings: value based threshold

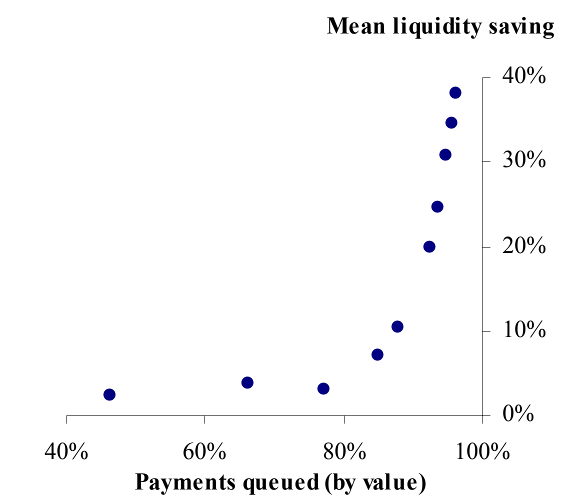

Chart 5: Non-linear profile of increase in average settlement time (AST): value based threshold

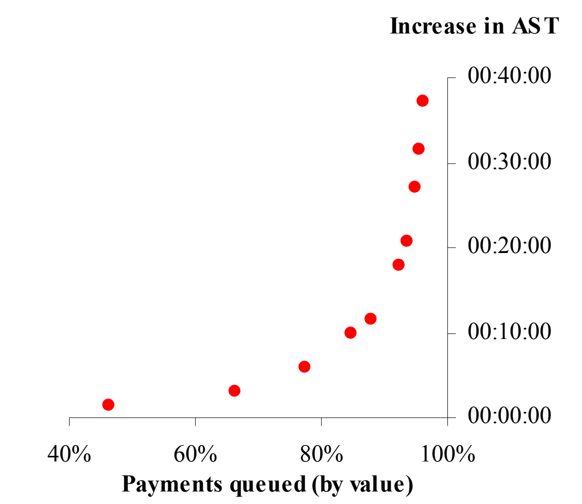

Charts 4 and 5 above illustrate the non-linear relationship between values queued and the liquidity saving and average settlement time metrics. The nature of this relationship suggests that the benefits of RRGS depend on having a critical mass of payments whose submission is being coordinated by the queue-release algorithm; for CHAPS this point is reached where 80-90% of payments by value are queued. It seems likely that the proportion of queued payments beyond which RRGS starts to offer significant savings will vary depending on the characteristics of the payment system under investigation, for example we might expect that where a system's aggregate recycling level under RTGS is significantly lower than that seen in CHAPS, the effect would emerge at lower proportions of payments. Testing this proposition would be a useful extension to our analysis.

Although our aggregate results show that mean liquidity savings are similar under the value and volume based time-criticality thresholds (Section 3.1 – Table A), we observe significant differences in the volatility of these savings across the two methods. The standard deviation *increases* when using the value based approach but *decreases* (although by less than the mean) under the volume based approach. The charts below may provide some insight into this difference. Chart 6 shows the distribution of time-critical payments through the day by aggregating all payments sent in half an hour intervals, and taking the mean across the month. As shown by the blue and red bars, the mean profiles for the value and volume based thresholds are very similar, which is consistent with our finding that mean liquidity requirements are not influenced by the method of selecting time-critical payments. In contrast, we see that the standard deviation of time-critical payments across the month is higher (and more volatile through the day) under the value based threshold. One likely factor here is that the number of payments classified as time-critical is much lower under the value based threshold than the volume based threshold. It seems plausible that this translates into a greater volatility in the liquidity demands under RRGS because it makes the timing and destination of payments which provide the liquidity to initiate the settlement of queued payments more unpredictable.

Chart 6: Value profile of time-critical payments in half hour intervals: *mean* across 19 days

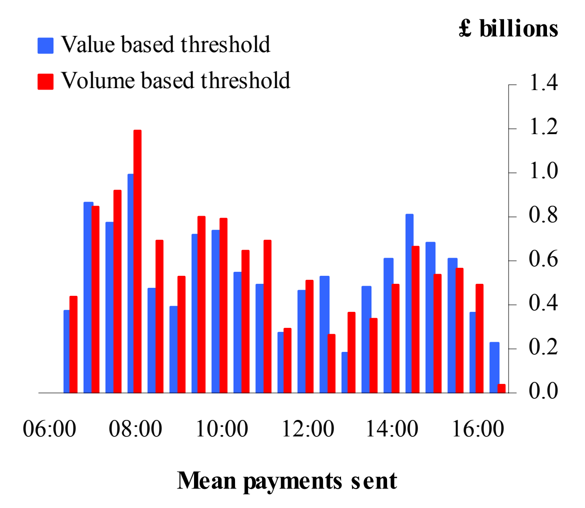

Chart 7: Value profile of time-critical payments in half hour intervals: *s.d.* across 19 days

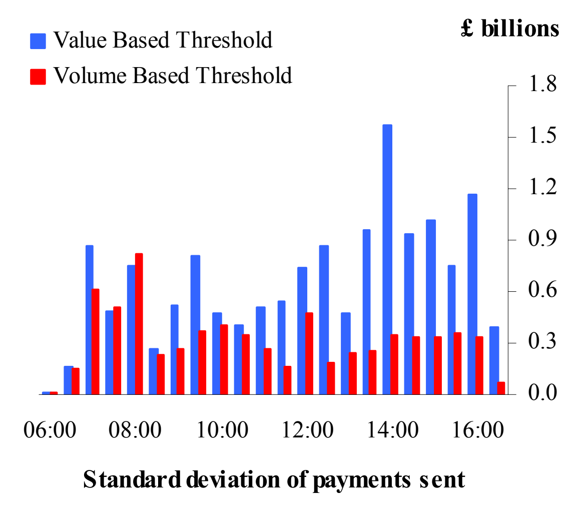

4.2 Distribution of RRGS benefits on individual banks

Table B in section 3.2 shows large variations in the impact of RRGS on individual banks. In this section, we explain the reasons for these variations.

Why does liquidity requirement under RRGS increase for some banks?

Charts 8-11 overleaf display the profile of payments sent (netted against incoming payments) before and after RRGS for two banks: A and B. Bank A's mean liquidity requirement slightly increases under RRGS, whereas Bank B is one of the biggest beneficiaries, seeing a large decrease in its mean liquidity requirement.

It appears that RRGS can potentially disrupt the recycling of payments for banks who already use liquidity very efficiently under RTGS, such as Bank A. Payments which they were due to receive from other banks, and subsequently use to fund their outgoing payments, are now being queued in the central scheduler (see the circled sections of Charts 8 and 9 below). Since they must make their time-critical payments without delay, they end up using more of their own liquidity to fund their outgoing payments under RRGS. In contrast, banks whose payment profiles are not as liquidity efficient under RTGS, such as Bank B, benefit the most from RRGS (Charts 10 and 11).

Chart 8: Profile of payments sent under <u>RTGS</u> (net of incoming payments) for Bank A: mean value at half an hour intervals across the month

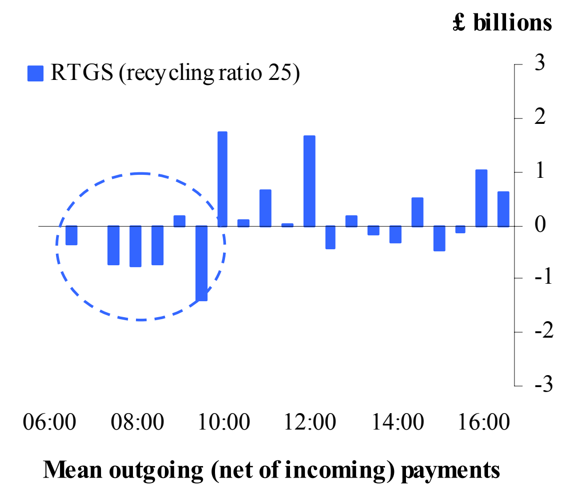

Chart 9: Profile of payments sent under <u>RRGS</u> (net of incoming payments) for Bank A: mean value at half an hour intervals across the month

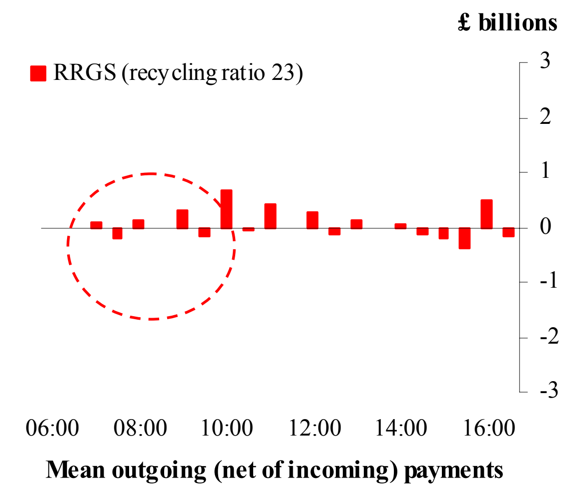

Chart 10: Profile of payments sent under <u>RTGS</u> (net of incoming payments) for Bank B: mean value at half an hour intervals across the month

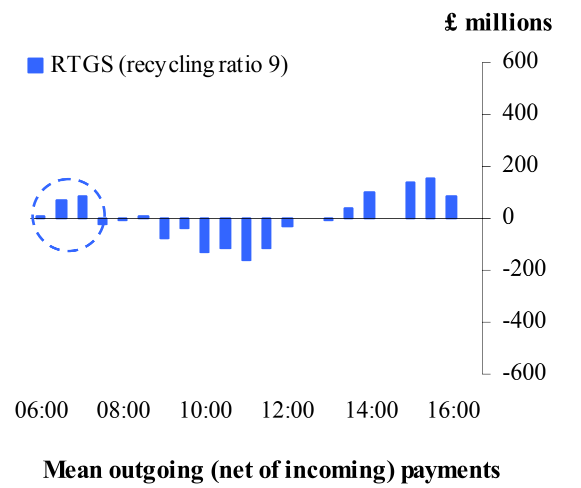

Chart 11: Profile of payments sent under <u>RRGS</u> (net of incoming payments) for Bank B: mean value at half an hour intervals across the month

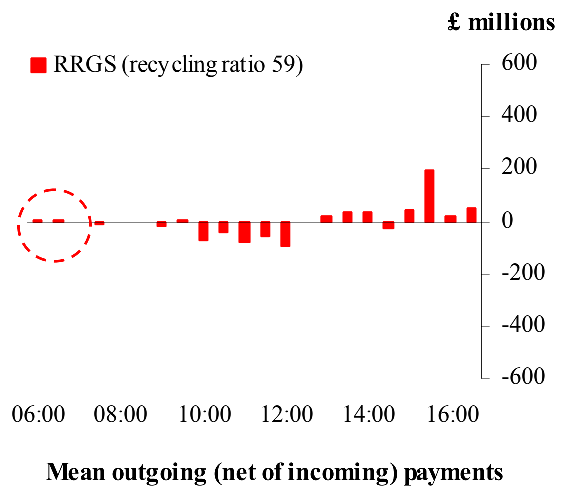

<u>Why do banks have different recycling ratios / how can RTGS be efficient for some of them?</u>

The answer lies partly in the intra-day profile of banks' payments. Becher *et al* (2007) suggest that structural differences in the underlying payment flows of banks may limit the extent to which payment timing can be managed so as to increase recycling ratios e.g. if certain banks (or their customers) routinely borrow in the overnight market, and others lend, the payment flows of the two groups will be correspondingly different. Their analysis of the intraday profile of net payments also reveals considerable variation, with some banks acting as net payers in the morning and net recipients later in the day, and other banks exhibiting the opposite.[19]

Charts 12 and 13 below also help shed some light on the issue. Chart 12 shows that banks with high recycling ratios under RTGS see their liquidity demands increase (although they still remain the users with the highest recycling under RRGS), whereas those with the lowest ratios see the greatest savings. Chart 13 shows that larger banks have better recycling ratios. This could be because a larger bank is likely to have more active links with the rest of the participants, increasing the probability of receiving incoming payments with which to recycle outgoing payments, especially in a system with extensive use of bilateral limits such as CHAPS. More fundamentally, the law of large numbers dictates that payment flows are more likely to be well balanced across the day where the volume of payments being made is greater.

Chart 12: Impact of RTGS recycling ratio on liquidity saving under RRGS
Payments ≥ £1bn time-critical

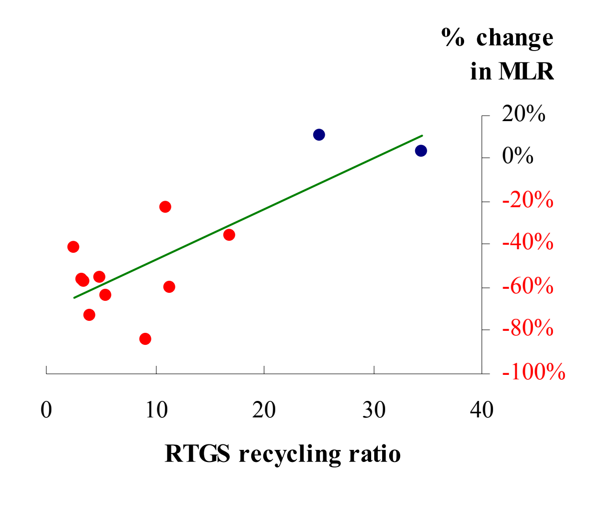

Chart 13: Relationship between bank size and recycling ratio under RTGS

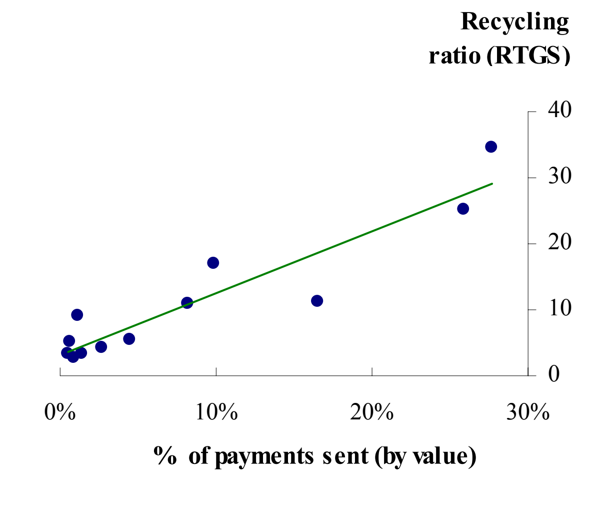

4.3 Possible biases due to the use of bilateral monitoring

One source of payment delay in CHAPS, as discussed in Becher *et al* (2007), is that CHAPS banks typically monitor the sending behaviour of other members on a bilateral basis, in order to enable them to quickly detect if either: (i) another member has suffered an operational problem which is preventing them from sending; or (ii) another member is deliberately delaying payments in order to reduce their liquidity needs. Where banks observe an interruption in normal flow of payments from a counterparty for either reason they will typically stop sending to that counterparty until normal flow is resumed. Incentives for bilateral monitoring will remain under RRGS, but we are not able to capture the effects of the constraints placed on payment submission due to bilateral monitoring in our simulation methodology. This implies that our methodology may underestimate the amount of settlement delay introduced by RRGS by unrealistically relaxing one set of constraints on payment submission. More generally, this highlights that

[19] We also intend to examine the profile of net payments of individual banks before and after RRGS to see whether this helps answer our question.

further work remains to be done carefully analysing the incentives created by the introduction of RRGS and how this would influence banks' payment sending behaviour.

4.4 Translating liquidity savings into cost savings

As discussed in Section 2, it is unrealistic to expect banks to post their *exact* daily minimum liquidity requirements into the system *ex ante*. Although the summary statistics we report are good indicators of the impact of RRGS functionality on banks' liquidity demands, their impact on banks' collateral posting decisions is not very clear. In some simulations, for example, RRGS reduces aggregate mean daily liquidity requirements across the month, but increases the aggregate of the maximum liquidity requirement across the entire month . So we need to assess how banks react to changes in the mean *and* standard deviation of their liquidity demands.

This is done by following the method employed in James & Willison (2004) of taking observed collateral posting decisions of CHAPS members and using regression analysis to estimate the influence of the mean and standard deviation of banks' max liquidity requirements on this decision. We use data for July-December 2006 to carry out a panel regression where the explanatory variables are the mean and the standard deviation of maximum collateral used on day t calculated over the 30 previous days in the sample; and the Libor / repo spread lagged by one day. In contrast to the methodology of James & Willison, a generalised least squares (GLS) estimator is used with correction made for serial correlation across time and within bank. The regression results are shown in Table D:

Table D: Cross-sectional time-series GLS regression of collateral posted: all banks
Including reserve account balances

	Coefficient	Standard error
Mean of maximum collateral used	0.24*	0.05
Standard deviation of maximum collateral used	0.98*	0.07
Libor / repo spread	0.00	0.00

*Coefficients marked with * are statistically significant at 1%*

All coefficients except the Libor / repo spread are in logs. The results could therefore be interpreted as: (i) a 1% increase in the mean liquidity requirement leads to a 24 basis point increase in collateral posting; and (ii) a 1% increase in the standard deviation leads to a 98 basis point increase in collateral posting.[20] The Libor / repo spread has a statistically insignificant impact on collateral posting, possibly because the true opportunity cost of posting collateral is low for banks subject to the UK Stock Liquidity Regime, who between them submit the majority of CHAPS payments by value.

We can now interpret our simulation results more clearly, by estimating the impact of RRGS on banks' collateral posting decisions using the above coefficients, our simulation results and the

[20] The collateral posting decision is four times more sensitive to the volatility in daily liquidity needs compared with the mean, suggesting that such decisions are not made on a daily basis. Banks may also choose to post collateral for precautionary reasons, for (unexpected) large payment outflows.

figure for mean collateral posting by a CHAPS bank in the 2nd half of 2006: £3.7 billion. This is shown in Table E below:

Table E: Estimating the impact of RRGS on banks' collateral posting decisions using James and Willison (2004)

Criteria for time-criticality	Early submission	% Δ Liquidity requirement		Effect of Δ in mean & standard dev on collateral posting	Settlement delay (hh:mm)
		Mean	St dev		
≥ £500mn	-	-10	+10	**+0.3bn**	**+00:12**
≥ £1bn	-	-38	0	**-£0.3bn**	**+00:37**
Random 10%	-	-12	-10	**-£0.5bn**	**+00:11**
Random 3%	-	-37	-18	**-£1.0bn**	**+00:25**
≥ £500mn	Random 50%	-15	+9	**+0.2bn**	**-01:20**
≥ £1bn	Random 50%	-37	+5	**-£0.1bn**	**-00:33**
Random 10%	Random 50%	-14	-12	**-£0.6bn**	**-01:15**
Random 3%	Random 50%	-39	-18	**-£1.0bn**	**-01:01**
≥ £1bn	Random 20%	-37	+2	**-£0.3bn**	**+00:05**
≥ £1bn	Random 80%	-37	+2	**-£0.3bn**	**-01:01**

The key result in the above table is that the increase in volatility (as seen by the increase in standard deviation) causes mean savings under some simulations to disappear, and for two simulations results in higher collateral postings for the system as a whole: banks can post less collateral because their mean liquidity requirement is reduced, but they must post more to cover the extra volatility. As discussed in Section 4.1, the increase in volatility is always observed with the value based time-criticality thresholds. Therefore under this metric the method of identifying time-critical payments does lead to significant differences in the costs faced by CHAPS banks.

4.5 Policy Implications

Chart 14 overleaf shows the results of simulations where the ≥ £1bn time-criticality threshold has been used. The *estimated* mean level of collateral that needs to be posted is plotted on the x-axis against the mean time of settlement on the y-axis, bringing together sections 4.3 and 4.4. Unsurprisingly, RTGS requires the highest level of collateral. The blue observation is an RRGS simulation where historical payment data is taken as exogenous (which may be unrealistic). This shows banks face a trade-off: a decrease in collateral costs coupled with an increase in settlement delay. Altering banks' payment submission behaviour is shown to be able to improve this trade-off, where a significant fraction of payments are submitted early we see unequivocal improvements in aggregate welfare compared to RTGS: estimated collateral costs are lower and mean time of settlement is earlier.

Chart 14: Estimated collateral posting vs settlement delay
Payments ≥ £1bn time-critical

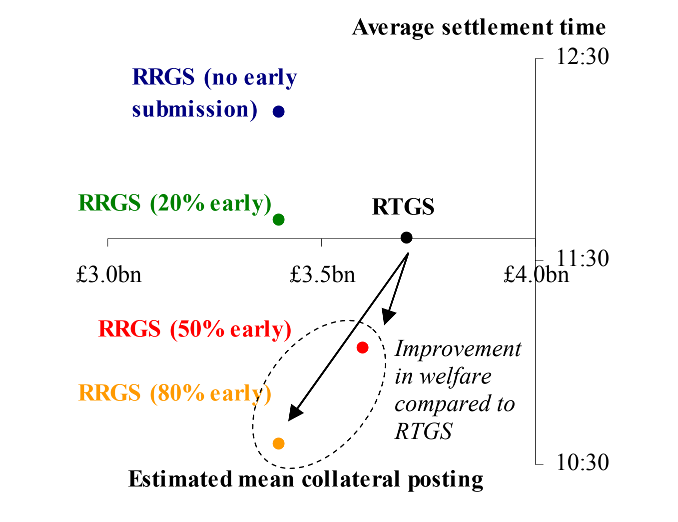

The results in Chart 14 give some weight to the idea that the introduction of hybrid functionality would be beneficial to CHAPS users. However, to make a firm policy recommendation about whether CHAPS should incorporate hybrid functionality further investigation would be needed to analyse the following:

(i) What is a realistic assumption about banks' level of knowledge of their outgoing payments at the start of the day?[21];
(ii) What would be the development costs of incorporating RRGS into CHAPS?;
(iii) Would banks trust the operational reliability of the central scheduler (anecdotal evidence suggests that this may be an issue)?;
(iv) What bias is introduced by our inability to incorporate the effect of bilateral monitoring on payment submission behaviour?;
(v) What opportunity costs do banks face both from posting liquidity in CHAPS and from settlement delays?

If after considering these factors a decision were taken that RRGS would be beneficial, there are additional queuing functionalities that might bring further benefits if introduced in parallel. For example, a *'settlement no later than'* option might encourage banks to submit even some *'semi time-critical'* payments to the central queue (i.e. payments which for example, have to be settled in two hours, but can be queued until then for potential offsetting). If the *settlement no later than* time is reached and a payment has not been released for RRGS settlement, the queue places the payment into the RTGS stream for immediate settlement.

Gridlock resolution might also provide further optimisation, where large payments hold up conditional release of smaller payments in the queue because of the FIFO rule. However, our analysis suggests that only around 2-3% of all payments remain in the queue at end of day and

[21] Our dialogue with CHAPS users suggests a figure in the region of 20-50% may be realistic, although with significant variations seen across banks, primarily depending on the proportion of their payments that are associated with financial market flows.

need to be settled multilaterally net. This is the case even where 96% of payments by value are submitted to the central queue and suggests that the FIFO rule works well for CHAPS.

Finally, the fact that savings from RRGS are unevenly distributed, and some CHAPS users may even face cost increases, suggests that some coordination problems may be faced in getting general agreement on investment in RRGS. However, the fact that the smallest CHAPS banks get most benefit from RRGS is of interest to policy makers. One area of interest for the Bank of England over recent years has been in analysing the exposures created in CHAPS due to the highly tiered nature of its membership. As Harrison *et al* (2005) explains, a tiered structure where a large number of users make payments across the system indirectly through a correspondent bank has the potential to create credit exposures between correspondent banks and their customers. Therefore, to the extent that RRGS functionality might make direct membership of CHAPS more attractive to smaller banks, this might be favourably regarded by policy makers.

5. Conclusions

Our simulations show that CHAPS banks could benefit from the introduction of receipt reactive functionality. Two distinct improvements could potentially be captured: a reduction in the amount of liquidity that banks would have to post to settle their payments, and earlier settlement of those payments. It seems that a key determinant of the magnitude of the liquidity saving would be the proportion of payments that could be submitted to the receipt-reactive queue, while the improvement in settlement timing would depend on the proportion of payments whose submission is currently delayed.

Our analysis of the distribution of benefits across banks shows that any liquidity savings achieved would not be evenly shared across banks. This reflects the fact that in CHAPS banks with fewer payments typically face a proportionally higher liquidity need to settle these payments under a RTGS design. It is these banks who would benefit most from the introduction of a RRGS design. By contrast a subset of large banks who currently achieve high recycling ratios would see no savings and may even face a small increase in their liquidity needs.

We show that the value profile of time-critical payments can impact on *collateral* savings achievable under RRGS, through their influence on the volatility of liquidity demands. Where it is the largest payments by value which are time-critical, collateral demands are typically higher than where values of time-critical payments are more evenly distributed.

Through analysing a synthetic payment dataset we confirm that a key determinant of the impact of RRGS functionality is the liquidity recycling ratio that banks achieve under RTGS. We believe that this in turn is influenced by a number of factors including the number of direct members of the payment system, the volume and value of payments being processed in that system, the network topology of the system and the profile of payment values. This may shed light on the observation that the majority of payment systems which have already adopted hybrid designs have historically always had a large number of direct participants.[22] Comparison of our results using the CHAPS and the synthetic datasets suggests that the existence of decentralised co-ordination mechanisms, such as bilateral monitoring, is already delivering some liquidity savings to CHAPS banks by reducing the volatility of their liquidity needs.

We have identified several promising avenues along which our work could be extended. When the Bank of Finland payment simulator is capable of replicating balance-reactive functionality then simulating the impact of its introduction to CHAPS would complement our analysis. A more detailed examination of the incentives banks face to submit payments to the RTGS stream would be worthwhile, especially to the extent it allows us to better understand what bias is introduced by our relaxation of bilateral constraints in our simulations. The differences in our findings depending on the method adopted to designate time-critical payments suggest there would be merit in exploring whether greater benefits could be achieved if banks co-ordinated on particular patterns of submission of time-critical payments early in the day to kick-start the release of queued payments. Finally, it would be interesting to explore the impact of RRGS under stressed circumstances i.e. where one or more CHAPS banks face an operational problem.

---

[22] E.g. BI-REL in Italy (120 in September 2004) and RTGS^{plus} in Germany (93 in December 2003): BIS (2005).

Annex 1: Methodology for generating synthetic payments dataset

The basic method adopted to generate our artificial payments data was to use a Matlab program which implemented the following procedure:

(i) Generate 4 columns (denoted A-D) of numbers of length n (the number of payments being generated) drawn from a log normal distribution. Generate an additional column (E) of length n drawn from a uniform distribution between 0 and 1.

(ii) Rescale and round to the nearest integer the numbers in A-C so that they contain integers between 1 and b, where b is the number of banks being simulated.

(iii) Rescale the numbers in D so that it contains numbers between 1 and m, where m is the largest possible payment being simulated. D represents the values of payments.

(iv) Compare matching rows of A and B. If $A_i = B_i$, *and* $E_i > 0.5$, *replace* B_i *with a number taken from C. Continue taking consecutive numbers from C until* $A_i \neq B_i$. *Follow the same procedure, but replacing* A_i *with a number from C, if* $E_i \leq 0.5$. *Repeat process until all* $A_i \neq B_i$. *A denotes the sending bank, B denotes the receiving bank, the procedure is done to ensure no payments are sent and received by the same bank. Draws from* E_i *to decide whether the sending or receiving bank is replaced are intended to ensure that no bias is introduced to the probabilities of a particular bank being a sender or receiver of payments.*

(v) Combine A, B and D with a column of payment times (T) which contains n entries evenly spaced between o (system open) and c (system close) to complete the dataset of artificial payments.

One extension to this method we used to make the dataset more realistic was to add some squaring-off payments at defined intervals through the day, which have the effect of ensuring that all banks end the day with zero balances. At the first two squaring-off points (after payments n/2 and 3n/4) each bank squares off half its accumulated balance, provided that balance has the same sign as its final balance at end of day. At end of day the remaining position is squared-off. Rather than explicitly modelling squaring-off payments between banks we assume that all squaring-off payments are made with Bank b+1 which is introduced solely for this purpose. Throughout we quote results from a dataset including squaring-off payments.

We believe the introduction of squaring-off makes our dataset more realistic. In practice payment banks typically face a non-zero net payment flow on behalf of customers on each day, as payments sent are unlikely to perfectly match payments received. However, this net redemption is typically managed through overnight market transactions with counterparts, leaving most banks with final net balances that are close to zero.[23]

---

[23] The residual size of net flows is to some extent dictated by the monetary policy implementation regime being followed in a country. For example in the UK, since the introduction of reserve averaging in May 2006, banks do not face a requirement to completely square-off their net positions as this can be absorbed to some extent by changes to their overnight balances at the Bank of England. By contrast, prior to the introduction of reserve averaging, a strong incentive existed for banks to square-off to zero.

References

Bank for International Settlements (1997), 'Real-time gross settlement systems', *Committee on Payment and Settlement Systems Publications No 22.*

Bank for International Settlements (2005), 'New developments in large-value payment systems', *Committee on Payment and Settlement Systems Publications No 67.*

Bech, M, Christine, P and Soramaki K (2007), 'Global trends in large value payments', *Mimeo*

Becher, C, Galbiati, M and Tudela, M (2007), 'The timing and funding of CHAPS Sterling payments', *Bank of England Working Paper (forthcoming)*

Harrison, S, Lasaosa, A and Tudela, M (2005), 'Tiering in UK payment systems: credit risk implications', *Bank of England Financial Stability Review December 2005, pages 63-72.*

James, K and Willison, M (2004), 'Collateral posting decisions in CHAPS Sterling', *Bank of England Financial Stability Review December 2004, pages 99-104.*

Johnson, K, McAndrews, J and Soramaki, K (2004), 'Economising on liquidity with deferred settlement mechanisms', *Federal Reserve Bank of New York Economic Policy Review, December 2004, pages 51-72.*

Leinonen, H and Soramaki, K (2003), 'Simulating interbank payment and securities settlement mechanisms with the BoF-PSS2 simulator', *Bank of Finland Discussion Papers 23/2003.*

Martin, A and McAndrews, J (2007), 'Liquidity-saving mechanisms', *Federal Reserve Bank of New York Staff Report no 282.*

McAndrews, J and Trundle, J (2001), 'New payment system designs: causes and consequences', *Bank of England Financial Stability Review, December 2001, pages 127-136.*

Millard, S and Polenghi, M (2004), 'The relationship between the overnight interbank unsecured loan market and the CHAPS Sterling system', *Bank of England Quarterly Bulletin, Q1 2004, pages 42-47.*

Willison, M (2004), 'Real-time gross settlement and hybrid payment systems: a comparison', *Bank of England Working Paper No 252.*